

REMARKS

Claims 1-29 are pending in the present application. Reconsideration of the claims is respectfully requested.

Applicants would initially like to acknowledge and thank the Examiner for taking the time to clarify the basis for the rejection of Claims 28 and 29. While no agreement was reached, the Examiner's explanation of the reasoning of the rejection of such claims providing Applicants with assistance in preparing this response.

I. 35 U.S.C. § 102, Anticipation

The Examiner rejected Claims 1-6, 11-18 and 23-29 under 35 U.S.C. § 102 as being anticipated by U.S. Patent 6,240,529 to Kato. This rejection is respectfully traversed.

The Claimed Invention

The claimed invention advances the state of the art by providing improved automation capabilities for the debugging of one or more processes in a multi-process environment. The improved debug automation advantageously provides capabilities previously unattainable by providing the ability to automatically save *and* automatically restore a process state. Such automation provides significant improvements in debugging by allowing a process to run until some event occurs, saving the process state into a checkpoint file, modifying registers or other variables, and then resuming execution - all under program control. This is particularly useful for long running programs which may not crash until several hours or days of running. For example, it is possible to automatically resume execution from a saved process state in response to a predefined event, providing improvements in debug capabilities by allowing a programmer who is debugging a process to instruct the debugger to run between a checkpoint and breakpoint repeatedly using a plurality of different register or memory variable values in order to assist in problem determination.

In particular, the present invention provides a debugger combined with a child process that can be checkpointed and restarted. This may be done automatically under program control of the debugger. In a preferred embodiment, the debugger, executing in memory as a process, creates a child process for the program being debugged. The child

process can then, in turn, create further child processes. The debugger, which is the parent process, has control over the child process(es). The debugger may save the image of all processes under control of the debugger in a checkpoint file to resume debugging from that state. Thus, the present invention provides a mechanism for automatically resuming debugging from a saved state, allowing the programmer to modify registers and memory variables and resume debugging from such known state (Specification page 9, lines 17-31). This use of a debugger to recreate processes from a checkpoint file can potentially save a user several hours or even days of debugging efforts to debug a long-running process that crashes after several hours or days of execution if a checkpoint file containing a process image is created on some interval-basis (Specification page 10, lines 19-30).

A key enabling feature of the present invention is the ability to resume processing from a known state (that was previously saved) under process control of a parent debug process. Since the debugger process creates or spawns the child process which is the subject of the debug session, the debugger process can be programmed to checkpoint and restore the child process using a multitude of different operating conditions such as variable or memory values (Specification page 11, 1-10), thus significantly advancing the state of the art.

Claim 1

With respect to Claim 1, such claim recites the claimed step of "retrieving the stored process state in response to a predefined event". This predefined event is a key enabling feature of the present invention, allowing for retrieval of a known process state in response to this predefined event to thereby allow subsequent execution of the process from a known state. In rejecting Claim 1, the Examiner cites Kato Col. 8, lines 29-35 as reading on this claimed step. Applicants show that to the contrary, this passage states:

"When a state restoration command is issued, storage situation management unit 121 displays a list of currently existing state storage files 118 based on storage situation management file 122. FIG. 4 shows an example of the list. If the file to be restored is selected from the

information, then state storage restoration unit 117 reads in the file and restores the debugged state.”

It can clearly be seen that Kato’s restoration is responsive to a state restoration command *currently* being issued. This does not teach or otherwise disclose that claimed step of retrieving the stored process state *in response to a predefined event*. The cited reference requires a manual restoration of a state storage file as selected by a user from a list of existing files (Kato Col. 4, lines 20-24; Col. 8, lines 29-35; Figure 4; Col. 10, lines 3-10; Col. 11, lines 17-29). Kato discusses problems with retrieving a state storage file at Col. 4, line 62 – Col. 5, line 11, but this retrieval is *responsive to a user input command* (Kato Col. 4, lines 20-24). In fact, the primary purpose of the Kato teachings is to assist the user in selecting which file from a plurality of saved state storage files to choose (Kato Col. 11, lines 17-28). This is accomplished by saving state storage files names in a trace buffer so that the user can look at the trace buffer to determine which state storage file had just been saved prior to a system crash (Col. 10, line 47 – Col. 11, line 16). The key point being that the teachings of the Kato reference are directed to assisting the user in manually selecting from one of many state storage files to be used when the user manually issues a state restore command (see, for example, Kato Figure 6, element 407). The teaching of a user entering a manual state restore command – albeit under assistance of state storage file names being included in a trace buffer to assist in such manual selection – cannot reasonably be construed to read on the claimed feature of retrieving the stored process state *in response to a predefined event*. The claimed feature of retrieving the stored process state in response to a predefined event advantageously allows for a process to automatically run repeatedly between a checkpoint and breakpoint using a plurality of register or memory variable values to thereby improve debugging capabilities (Specification page 16, line 28 – page 17, line 3). As the cited reference does not teach such retrieval in response to a *predefined* event, but rather responsive to a *current* input command, the cited reference does not anticipate Claim 1.

Nor would it be obvious to modify Kato to include such a retrieval of state information in response to a predefined event. Kato teaches that state information can be *stored* when a certain event occurs (Col. 9, lines 64-67), but requires manual user

intervention in the *restore* of such state information (Col. 10, lines 3-10; Col. 11, lines 17-29). Since Kato was aware of a desire to automatically store state information, but yet provided no solution for retrieval of state information in response to a predefined event, it must not have been obvious to Kato as to how to accomplish this. In addition, Kato's express purpose to improve retrieval of state information is to simultaneously store situation information in a storage situation management file when state information is stored, to facilitate a *manual determination* of which among a plurality of existing state storage files should be manually selected by the user (Col. 5, lines 35-48) based upon the user's manual input of a restore command at the time the restore command is desired (Col. 10, lines 3-10; Figure 5B, element 322). Modifying Kato in accordance with the claimed invention would defeat this expressed purpose.

Claim 6

With respect to Claim 6, the cited reference does not teach the claimed step of "reinitiating debugging from the stored process state, wherein the process has control over at least one child process and the process state includes a process descriptor for each of the at least one child process". Notably absent from the Examiner's reasoning in rejecting Claim 6 is any mention or discussion of this claimed step. The Examiner merely asserts that the cited reference teaches "reinitiating debugging from the stored process state". Applicants show that Claim 6 goes beyond this assertion, and specifically recites "wherein the process has control over at least one child process and the process state includes a process descriptor for each of the at least one child process" in addition to the alleged "reinitiating debugging from the stored process state". For a prior art reference to anticipate in terms of 35 U.S.C. 102, every element of the claimed invention must be identically shown in a single reference. *In re Bond*, 910 F.2d 831, 15 USPQ2d 1566 (Fed. Cir. 1990). As every element of the claimed invention is not identically shown in a single reference – and in fact has not even been alleged to be shown in a single reference – Claim 6 is shown to have been erroneously rejected under 35 U.S.C. § 102 as being anticipated by U.S. Patent 6,240,529 to Kato.

This invention of Claim 6 advantageously enables debugging in a multi-processing environment. In a traditional multi-processing environment, a process

typically creates or spawns child processes which run concurrently with the parent process (Specification page 9, lines 7-16). In order to provide debugging capabilities in such a multi-processing environment, per the claimed invention, it is critical to keep track of these child processes *in addition to* the process being debugged. Inclusion of process descriptors for each child process in the process state of the process being debugged provides this capability. The cited reference only alludes to maintaining debug information for a single running process. Thus, the invention of Claim 6 provides a significant advancement in the art by providing debug capabilities in an environment with multiple concurrently running processes.

Claims 26 and 27

Claim 26 depends upon Claim 1 and is shown to not be anticipated by the cited reference for reasons given above regarding Claim 1.

Further with respect to Claim 26 (and dependent Claim 27), such claim recites that the predefined event (to which the retrieval of the stored process state of Claim 1 is responsive to) is a checkpoint. This checkpoint is shown in the preferred embodiment by restore checkpoint element 616 in Figure 6. As described at Specification page 16, line 28 – page 17, line 3, by providing the ability to restore an image of the process being debugged in response to an event – and in this case a restore checkpoint – the debugger may advantageously be instructed to automatically run between a checkpoint and a breakpoint repeatedly. This claimed feature advantageously provides the ability to repeatedly run between the checkpoint and breakpoint using a plurality of register or memory variable values to assist in debugging the process.

The cited reference does not teach the claimed steps of “retrieving the stored process state in response to a predefined event” ... “wherein the predefined event is a checkpoint”. As can clearly be seen by this claim language, the step of retrieving the stored process state is in response to a checkpoint. The Examiner cites Kato Fig. 5A, #309 and #314 and Col. 9, lines 15-36 as reading on this claimed step. Applicants show that to the contrary, step #309 merely checks whether an event has occurred and if so, processing of the action registered in connection with the event is performed (Kato Col.

9, lines 15-21). There is no mention whatsoever of any type of specific action of *retrieval of a stored process state in response to this event checking*.

Nor does Kato's teaching with regard to step #314 overcome this deficiency. Step #314 determines if a break occurred. If YES, the execution flag is changed to OFF and execution processing ends (Col. 9, lines 35-38). If NO, a determination is made on *whether to continue execution of instructions* (Col. 9, lines 38-43). Again, this has absolutely nothing whatsoever to do with retrieval of the stored process state in response to a checkpoint, as claimed in Claim 26. Therefore, Claim 26 (and dependent Claim 27) is shown to not be anticipated by the cited reference as there are missing claimed elements.

Further with respect to Claim 26 (and dependent Claim 27), Applicants show that the cited reference does not teach the claimed step of "repeatedly running between the checkpoint and the breakpoint for a plurality of times". The Examiner cites this same passage at Kato Col. 9, lines 15-36 as teaching this claimed feature. This passage is repeated below in its entirety to clearly show there is no mention whatsoever of repeatedly running between the checkpoint and the breakpoint for a plurality of times:

"Then, it is checked whether or not a state corresponding to one of events registered by the user is satisfied as a result of the execution of the instruction (309), and if an event is detected (310), then processing of an action registered in connection with the event is performed (311).

By the way, in the present embodiment, a function to storage a debugged state into a file if a certain event occurs is provided in addition to the conventional functions. Thus, it is assumed that an event for which storage processing of a debugged state is to be performed when the event occurs is registered in advance. In this instance, if the event is detected, then a state storage request is generated (312), and a current debugged state is stored into state storage file 118.

In the present embodiment, since also the function to manage a situation when a debugged state is stored is additionally provided, the situation upon the storage is registered simultaneously into storage

situation management file 122 (313). If a request for a break is generated in response to the detection of the event (314), then the execution flag is changed to OFF (316), thereby ending the execution processing.”

Applicants show that this cited passage has nothing to do with “repeatedly running between the checkpoint and the breakpoint for a plurality of times”. Rather, it teaches that if an event is detected, processing of an action registered in connection with the event is performed (Col. 9, lines 15-20). This is a general statement that makes no mention of what the registered action is or does. It merely states that the registered action ‘is performed’. The cited passage goes on to state that a state storage request is generated and the current debugged state is stored into state storage file 118 (Col. 9, lines 27-29) along with the situation (Col. 9, lines 30-35). This is not seen to teach or otherwise disclose in any fashion the claimed step of “repeatedly running between the checkpoint and the breakpoint for a plurality of times”. Hence, Claim 26 (and dependent Claim 27) is further shown to not be anticipated by the cited reference.

Claim 28

Claim 28 advantageously provides a method for debugging a process in a multi-process environment. This is accomplished by initiating a debug process and creating a child process from the debug process. As described above in the Summary of the Invention section, creation of the child process from the debug process advantageously allows the debug process to control the child process (which is the primary focus of the debug session due to the recitation in the claim of saving a process state of the child process, retrieving the stored process state, and executing the child process using the stored process state) and thus advantageously provides an ability to completely automate checkpoint and restore operations for the child process. In one example of such automation enablement, it is possible to program the debug process to perform multiple checkpoint and restore operations of the process it is controlling (the child process) while varying various other system parameters such as registers and memory to thereby facilitate a more robust debug session.

With respect to Claim 28, Applicants show that the cited reference does not teach the claimed steps of “creating a child process from the debug process”, “saving a process state of the child process” (the child process having been created from a debug process) and “executing the child process using the stored process state”. In rejecting Claim 28, the Examiner states that Kato Col. 7, line 15 reads on the claimed step of initiating a debug process, and that Kato Col. 10, lines 43-46 reads on the claimed step of creating a child process from the debug process. Applicants show error in this assertion as follows.

Claim 28 recites two processes – a debug process and a child process (which is created from the debug process). Col. 7, line 15 of Kato (which the Examiner asserts reads on the claimed debug process) reads as follows:

“When debugging is started a down load command is inputted from information IN/Out apparatus 102 to down load object code file 104 of a program *which is an object of debugging to the debugging apparatus.*”
(emphasis added by Applicants)

So, according to this reading of Kato, the debug process is the object of debugging. Of particular importance is that according to Claim 28, the *child process* (created from the debug process) is the object of debugging, as it goes on to state “saving a process state of the child process”, retrieving the stored process state” and executing the child process using the stored process state”. Per the Examiner’s interpretation of Kato, the *debug process* is the object of debugging (as the Examiner states that Kato Col. 7, line 15 reads on the claimed debug process). This is not merely a play on words, but a very real difference in that the claim recites that the process being debugged (the child process) is created from another process (the debug process). This claimed feature advantageously allows the debug process to control and thus automate debugging of the child process. As the passage cited by the Examiner states it is the debug process which is the object of debugging, which is contrary to what is claimed (the child process created from the debug process is the object of the debugging), it is shown that Claim 28 is not anticipated by this reading of Kato.

Nor does the Kato passage cited at Col. 10, lines 43-46 overcome this deficiency. There it states (the entire sentence is being reproduced herewith):

“When setting so that trace information is also acquired simultaneously at a timing of storage of a debugged state is performed in this manner, in Embodiment 2, the storage situation management unit in Embodiment 1 additionally has a function of recording a debugged state storage timing into trace information.”

This passage states that debugged state storage timing is stored into trace information. This merely states that two things are being stored, not that two processes are running, where one process (the child process) is created by the other and this one process (child process) is the process being debugged (by saving and retrieving its associated state information). Restated, a teaching of storing two things in a file does not teach or otherwise disclose a debug process that creates a child process, where the child process is the object of debugging (by saving and retrieving the child's process state for use during subsequent execution of the child process).

The Examiner goes on to discuss various aspects of Kato as reading on the claimed saving, retrieving, and executing steps. These are all with respect to the object code file 104 which is the object of debugging (Kato Col. 7, lines 11-20), which the Examiner has equated with the claimed debug process. As stated above, the debug process and the child process are very different, with a unique relationship where the child process (the subject of the claimed steps of saving, retrieving, and executing) is created from the debug process. Thus, it is shown that saving, retrieving and executing of a debug process is very different from the claimed steps of saving, retrieving and executing from a child process created from the debug process. Therefore, Claim 28 is further shown to not be anticipated by the cited reference.

Claim 29

Claim 29 recites saving two different process states – the process state of a traced process and the process state of another process that is not being traced. The claimed

feature advantageously allows checkpointing processes and saving their process images in a checkpoint file even when they are not being traced (Specification page 10, lines 19-21).

In particular, Claim 29 recites "saving a process state of a traced process and a process state of another process that is not being traced by the debugger". The cited reference does not teach or otherwise disclose saving a process state of a traced process *and* a process state of another process that is not being traced. The Examiner cites Kato Col. 7, lines 35-47 and 65-66 as reading on this claimed step. These passages are reproduced below:

"Debugger unit 106 in debugging apparatus 103 performs processing regarding a debugging function for breaking, tracing or the like. A break point at which execution of a program is interrupted when a certain event occurs and a trace wherein, when a certain event occurs, an instruction execution situation, a memory access situation and so forth at the point of time are recorded and individually set with commands in advance. The events set with the commands and actions such as breaking and tracing when such events occur are registered in event processing unit 110 and action processing unit 111 in debugger unit 106 and are processed in the following manner when the instructions are executed." (Col. 7, lines 35-47)

"In the debugging apparatus of the present embodiment, event processing unit 110 has, in addition to the conventional functions, a function of storing a debugged state into a file if a certain event occurs." (the entire sentence encompassing Col. 7, lines 65-66)

As can be seen, the passage at Col. 7, lines 35-47 generally discusses a debugging apparatus that has the ability to set of a breakpoint and perform tracing when a certain event occurs, such as an instruction execution situation or a memory access situation. These are well known debugging operations and are not seen to be of importance with

regards to Claim 29 and specifically the saving of process states from two different processes. The passage cited at Col. 7, lines 65-66 merely describes an ability to store a debugged state into a file if a certain event occurs. Thus, it is shown that this passage merely recites storing a debugged state if a certain event occurs. It does not teach or otherwise disclosure saving a process state *of the traced process and* a process state of *another process that is not being traced by the debugger*. At best, it teaches storing a single process state, whereas the claim is directed to storing two different process states from two different processes – with one state being traced and the other one not being traced. Thus, it is shown that Claim 29 is not anticipated by the cited reference.

Other Claims

Applicants traverse the rejection of Claims 2-5, 11 and 12 for reasons given above regarding Claim 1, of which Claims 2-5, 11 and 12 ultimately depend upon.

Applicants traverse the rejection of Claims 13-17 for similar reasons to those given above regarding Claim 1.

Applicants traverse the rejection of Claim 18 for similar reasons to those given above regarding Claim 6.

Applicants traverse the rejection of Claims 23 and 24 for similar reasons to those given above regarding Claim 13 and Claim 1.

Applicants traverse the rejection of Claim 25 for similar reasons to those given above regarding Claim 1.

Therefore, the rejection of Claims 1-6, 11-18 and 23-29 under 35 U.S.C. § 102 has been overcome.

II. 35 U.S.C. § 103, Obviousness

A. The Examiner rejected Claims 7-9 and 19-21 under 35 U.S.C. § 103 as being unpatentable over U.S. Patent 6,240,529 to Kato, and further in view of U.S. Patent 5,560,009 to Lenkov et al. This rejection is respectfully traversed for similar reasons to those given above regarding Claims 1 and 13, of which these claims ultimately depend upon.

B. The Examiner rejected Claims 10 and 22 under 35 U.S.C. § 103 as being unpatentable over U.S. Patent 6,240,529 to Kato, and further in view of U.S. Patent 6,412,106 to Leask et al. This rejection is respectfully traversed for similar reasons to those given above regarding Claim 1, of which these claims ultimately depend upon.

Therefore, the rejection of Claims 7-10 and 19-22 under 35 U.S.C. § 103 has been overcome.

III. Conclusion

It is respectfully urged that the subject application is patentable over the cited references and is now in condition for allowance.

The Examiner is invited to call the undersigned at the below-listed telephone number if in the opinion of the Examiner such a telephone conference would expedite or aid the prosecution and examination of this application.

DATE: 9/3/03

Respectfully submitted,

Wayne P. Bailey
Duke W. Yee
Reg. No. 34,285

Wayne P. Bailey
Reg. No. 34,289
Carstens, Yee & Cahoon, LLP
P.O. Box 802334
Dallas, TX 75380
(972) 367-2001
Attorneys for Applicants

RECEIVED
CENTRAL FAX CENTER
SEP 03 2003

OFFICIAL